

# A system to restructure hypertext networks into valid user models

**Johan Bollen** and **Francis Heylighen**

*Vrije Universiteit Brussel, Leo Apostel Center; Krijgskundestraat 33, 1160 Brussels, Belgium.*

*(e-mail: jbollen @vub.ac. be)*

---

We have implemented an experimental system that automatically restructures hypertext networks according to their users' browsing behavior. The system applies link weights to the hyperlinks in the networks and updates these link weights according to three learning rules. The learning rules are based on how often a particular hyperlink is being traversed and operate on strictly local information of link traversals. Changes in network structure are fed back to users by dynamic link ordering according to descending link weight. The system has been shown to be able to structure random hypertext networks into valid representations of their users' browsing preferences in two WWW experiments and a simulation using a mathematical model of user navigation.

## 1. Introduction.

### *1.1. Information retrieval in the WWW.*

**T**he Internet and its associated WWW (World Wide Web) hypertext protocol <sup>2</sup> have been experiencing an exponential growth during the past few years. Not only the number of users and Internet servers but also the amount of electronic information stored has been growing at an astounding pace.

In spite of its popularity among publishers and users, the WWW doesn't seem to be entirely living up to expectations. With a growing number of pages and an increasing number of links, users are experiencing more and more difficulties in retrieving the information they require <sup>10, 27</sup>. The retrieval of information from (electronic) databases and other information repositories has been the domain of the science of Information Retrieval (IR). Techniques and methods developed in IR have focused on indexing methods, automated retrieval by queries, query relevance improvement, database management, user interface, document clustering and document similarity, etc. Certain IR techniques have been applied to the WWW. Query based search sites such as Lycos (<http://www.lycos.be/>) and AltaVista (<http://www.altavista.com/>) and indexing sites such as Yahoo (<http://www.yahoo.com/>) have been highly successful and are being used by millions of WWW user to locate and retrieve the information they are looking for. These query based search services however disregard the WWW's actual hypertext structure and treat it as an unconnected repository of electronic documents.

The WWW has been specifically designed with the intention of not only providing a means of inexpensive global storage of information but also of effortlessly linking this information in a meaningful manner. This characteristic of the WWW enables users to actively browse the network's structure and locate information by following the connections among its documents. WWW users can freely interact with a hypertext system without explicitly formulating or knowing their retrieval goals in advance. Retrieval goals can serendipitously take shape while browsing, depending on their previous browsing history, context and changing interests. In spite of its advances, traditional IR techniques seem to fall short of dealing with the interactive way in which information retrieval in the WWW takes place<sup>7</sup>. Worse, although the WWW and hypertext in general seem to offer users more freedom and a more intuitive means of retrieval, this does not necessarily mean they are superior to other more traditional media in terms of retrieval efficiency and ease of use<sup>9</sup>.

Improving the efficiency of human navigation and retrieval from hypertext networks such as the WWW seems to be a critical issue<sup>25</sup>.

### *1.2. Related work in user modeling.*

A number of systems have concentrated on alleviating the difficulties users experience in browsing hypertext networks and the WWW by better informing users about underlying network structure<sup>34</sup> by means of visualization, graphical maps, guided tours, etc. A number of learning effects might also be induced to help the user cope with specific hypertext systems,<sup>8,18</sup>. The problems concerning retrieval from the WWW can, however, not solely be attributed to incomplete user knowledge. A very complex interaction between factors such as the designer's preferences, the user's goal and interests, the specific interfaces, etc. seems to be involved.

One very important factor to consider is the WWW's static and often inadequate hypertext structure<sup>37, 39</sup>. WWW sites are implemented and linked by local designers who use their best insights and personal interests to order and link the information they want to publish. WWW designers thus to a large extent define the static context within which the human or automated retrieval of information will take place. A complete lack of information on the plans, goals and interests of the designers and the consequent semantic significance of hyperlinks in the network, hampers the implementation of more advanced, semantically driven applications for the WWW and hypertext in general<sup>19</sup>.

Since the static, semantic aspects of the hypertext linking eludes objective measures and possible applications to information retrieval, researchers have tried to approach the hypertext retrieval problem from a different angle. Modeling and measuring the characteristics of the user(s)

rather than handling the structure of hypertext networks could enable specific systems to filter, tailor and enhance the content and structure of existing hypertext networks to better fit the user(s) interests and characteristics. A large part of the literature on adaptive hypertext and hypermedia has been concerned with this problem <sup>4</sup>. Many authors have described ways to measure the user's preferences, interests or abilities, and to change the behavior of the hypertext system accordingly <sup>18</sup>. A typical system will for example query users for their educational level and change its responses in accordance with a consequent categorization of the individual user as an expert or layperson. Other systems implicitly infer user characteristics from, for example, the personal interests users express in reading certain pages <sup>35</sup>, to make the system preferentially present pages that correspond to these interests. Some systems even learn to categorize users and dynamically adapt the system's output. This approach has yielded some very interesting systems and results <sup>26</sup>. User modeling techniques have found widespread applications in systems for retrieval suggestions, link filtering, etc. Rather than changing the underlying information structure, these systems enhance human navigation in hypertext by filtering out irrelevant material and advise links and other services.

The WWW has a number of specific characteristics that set it apart from other hypertext networks such as its highly distributed nature and enormous breadth of content and semantic domain. Most interesting for the WWW are therefore those applications that do not have any preformed model of the user or the information domain, but rather construct their models dynamically from user behavior and relevance feedback.

Collaborative filtering techniques for example have focussed on how users' interests overlap and how this partial overlap can be used for recommendation systems such as Pattie Maes's Firefly (<http://www.firefly.com>) and Luis Rocha's TalkMine system <sup>30</sup>. The rationale for these systems is that if certain groups of users have similar profiles, they will also share similar interests. A group's interests could therefore be used to suggest documents or services to other, similar users. Generally one could say collaborative filtering systems construct dynamic models of group similarity and dissimilarity, and indirectly apply that model to estimate the relevance of documents for individual users. These systems have, however, not been used to support interactive information strategies like browsing which is common to information seeking on the WWW. They also require users to extensively and explicitly specify their interests and characteristics. The contribution of these systems is, however, that they can combine the know-how of large groups of users to aid individual users in retrieving the most relevant material.

Other systems for information retrieval are specifically aimed at helping users navigate hypertext networks, but rely on more direct measurements of document relevance and user characteristics.

The Letizia system <sup>23</sup> for instance tracks an individual user's browsing behavior, deduces a model of the user's interests and can, based on this information, recommend possibly interesting documents. The system uses a simple list of keywords as a model of the user's interest. It is interesting that the system uses past navigational choices to implicitly construct its user model. Other approaches like the Webwatcher project <sup>16</sup> accompany the user while browsing and mark links that may possibly be worthwhile to pursue. The system learns link relevance by annotating hyperlinks with users' keywords and uses reinforcement learning based on retrieval rewards. Webwatcher has, unlike many other systems, been experimentally tested for its efficiency in advising appropriate hyperlinks. Results show that its advice (using five combined recommendation methods) overlaps in 48.9% of all cases with actual user selections. Remarkably, within the same experiment recommendations based on link popularity produced the only marginally lower score of 41.9%.

Some systems like Alexa (<http://www.alexa.com/>) seem to combine the advantages of collaborative filtering and browsing support. Alexa uses a web browser (Netscape or Microsoft's Internet Explorer) plug-in to feed back user navigation to its main server. The browser plug-in relies on the main server to recommend possibly relevant connections to the user. The system has proven to be quite popular.

Some systems have focussed on how the structure of hypertext systems can be adapted themselves. Kaplan et al. <sup>20</sup> present the **HyperFlex** system that uses associative matrices to store information about user preferences. Values in the associative matrices are updated continuously when users choose different ordering of advised links. The system allows user profiles to be merged to combine different users' knowledge. Unfortunately it is not designed for the WWW because it **deconstructs** hypertext into separate text and hyperlinks and requires a specialized hypertext application using menu bars and panels.

### *1.3. Adaptive hypertext structure.*

These systems all share the characteristic that they regard hypertext networks as given, static structures whose shortcomings can only be circumvented by either ignoring or **deconstructing** the system's hyperstructure, or helping the human navigator cope with the existent structure as well as possible. This consideration led us to the development of our experimental Adaptive Hypertext system. Rather than implementing an intermediate system that functions as a third party

to alleviate the shortcomings of an existing hypertext network, our approach was to construct a system that updated the hypertext network structure itself. The system's aim is to structure the network into a collective user model: a representation of any group of users' preferred pattern of linkage. It will thereby combine the previously mentioned advantages of collaborative filtering techniques (group based recommendations) with that of the navigation aid approach.

The presented method depends on the fact that people in general have stable and preformed ideas about the associations of concepts<sup>28, 22, 36</sup>. The study of association norms<sup>36</sup> and the measurement of the meaning of concepts has demonstrated that these ideas overlap among groups of individuals and that they can be measured reliably<sup>28, 24</sup> over large groups of subjects. Designers and users of hypertext networks also share overlapping notions of associativity and use these to respectively determine how the network's hypertext pages should be linked and organized, or how they should be navigated. When both user and design model on the other hand overlap poorly<sup>1</sup> the interaction between the user and the hypertext system (or any other tool) can become erroneous and inefficient<sup>29</sup>. The measurement of user and design models can thus be considered an important issue for the improvement of the interaction between users and (hypertext) systems. It is strongly related to the psychological measurement of meaning and word association. Most techniques in the psychological literature, however, rely on procedures in which large groups of human subjects explicitly name or indicate the associations to a certain word<sup>22</sup> and therefor have only limited applicability to the WWW due to its open-ended structure and content. Models of the users' overlapping notions of associativity (also referred to as mental models) could, however, be derived from implicit measures of general browsing behavior. The fact that groups of users consistently select a hyperlink between two hypertext documents could for **example** be used as an indication that the two hypertext pages are associatively related. The overlapping choices of a sufficiently large group of users to use certain hyperlinks and not others can thus be used as an implicit measurement of these users' notions of how hypertext pages should be linked. In analogy to the explicit measurement of word association norms one could say that the number of times a hyperlink on a page has been chosen compared to the total number of selections from that specific page will indicate the probability that a user visiting that specific page will select that hyperlink.

## **2. A network based feedback system for restructuring hypertext.**

### **2.1. Basic principles**

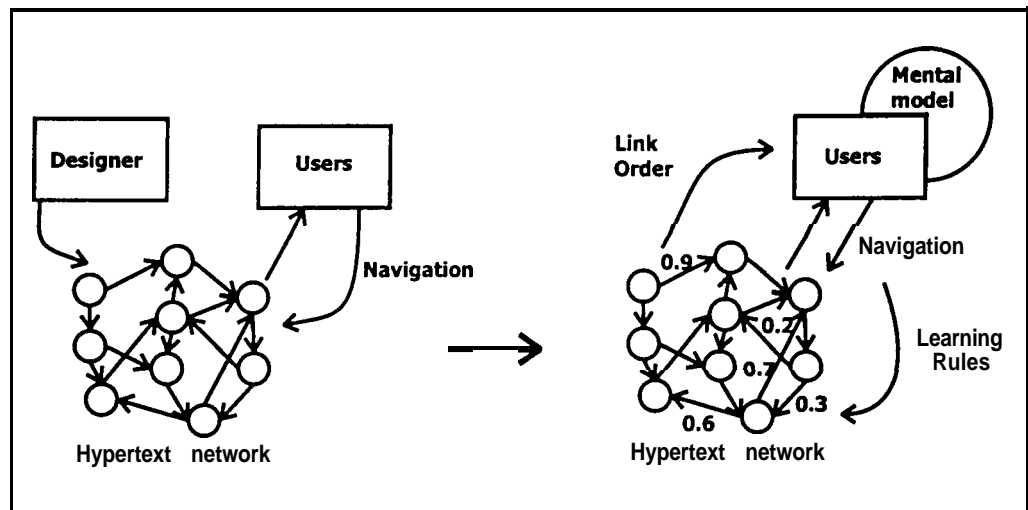
We have implemented a scheme for the dynamic restructuring of



hypertext that makes hypertext systems restructure themselves while they are being browsed, with little or no intervention from the human designer. The proposed system closes the feedback loop between users and the network by implementing the following (fig. 1):

- a) hyperlinks are assigned connection weights
- b) a set of learning rules changes connection weights according to the user's navigational decisions
- c) changes in network structure are fed back to the user by ordered presentation of links according to decreasing connection weight

Following is a short outline of the proposed system's main functionality. The implementation details will be discussed in section 3.1.



**FIG. 1:** The adaptive hypertext system closes the feedback loop between designer and user.

## 2.2. Weighted hyperlinks.

Hyperlinks are directional and Boolean; i.e. they point from one page to another and are either present or not present. They can not be modulated in terms of link relevance or quality. A certain form of link modulation can however be very useful. Human users need to assign strength of relation to the links from a given hypertext pages to be able to browse the network. Otherwise, all connections would be considered equally relevant and the user could not selectively navigate towards any specific position in the network. In order to enable the automatic restructuring of hypertext networks as intended by our system, we decided to enable the system to assign connection weights to hyperlinks. This would allow the system to modulate existing connections without their actual removal/creation and use the weights as evaluations of relevance.

## 2.3. Learning rules

The following set of three learning rules was implemented. They locally



change **connection** strengths according to whether users select a specific connection or not (fig. 2).

- 1) Frequency: This learning rule's function is analogous to Hebb's <sup>12</sup> model of human learning which states that if two concepts are simultaneously activated (or temporally close), the connection between these two concepts is reinforced. This principle of reinforcement lies at the heart of many models and systems for automated learning <sup>13, 21, 31</sup>. In analogy to Hebb's law of human learning the Frequency learning rule will reward the connection between two hypertext pages that has been traversed. Consequently, the more frequently a given connection is being used, the higher its connection strength will be.
- 2) Transitivity: The transitivity learning rules does not reinforce actually traversed connections, but rather introduces new and plausible connections to the network. When a user navigates from a certain node *a* towards a certain node *b*, and consequently navigates from node *b* to another node *c*, the transitivity learning rule reinforces the connection between node *a* and *c*. The transitivity learning rules thus tries to shorten retrieval paths by bridging plausibly related nodes. Whenever Transitivity introduces a new connection, however, this connection can only succeed in achieving sufficient connection strength if users think it is worthwhile and start using it. This usage in turn will lead to reinforcement by the Frequency learning rule. If users however feel the new connection is not relevant they will not select it and it will remain at its small, initial reward administered by the Transitivity learning rule. When, for example, a user browses a path through the nodes "school", "book" and "math", the transitivity learning rule will reinforce the connection between "school" and "math". After its introduction by the Transitivity learning rule, users will probably feel this connection is useful and will **therefor** start using it. The "school"-"math" connection will thus be reinforced by the Frequency learning rule. If, on the other hand, the user would browse a path through "school", "book" and "story", the transitivity learning rule will reinforce the connection between "school" and "story" but users will less probably feel this connection is useful and might choose to ignore it. It will consequently not be reinforced by the Frequency learning rule.
- 3) Symmetry: Connections in hypertext networks are directional and thus not necessarily symmetric (a certain hypertext page *a* can refer to another page *b* while vice versa this is not the case). It is nevertheless plausible that at least a degenerate form of symmetry holds for hypertext networks since they are associative networks. The Symmetry learning rule therefore enforces any connection from a

node **b** to another node a, whenever the connection between node **a** and node **b** has been traversed. Once a connection has been reinforced by Symmetry, users can either use the newly established connection while browsing or not. If they do select the connection, it will consequently be reinforced by Frequency. When for example a user browses from the node “word” to “book”, then the Symmetry learning rule will reinforce the connection between “book” and “word”. Users might find this connection useful because the two concepts are symmetrically related and its use will lead to reinforcement by the Frequency learning rule. If on the other hand, a user browses from “car” to “key”, Symmetry will also reinforce the “key” to “car” connection which users might find less useful or plausible. In this case, the Symmetry learning rule would have reinforced a “spurious” or at least less useful connection from the browsing users’ point of view and it will consequently be used less and receive less reinforcement from the Frequency learning rule.

The three learning rules operate strictly locally, i.e. during browsing, and in parallel

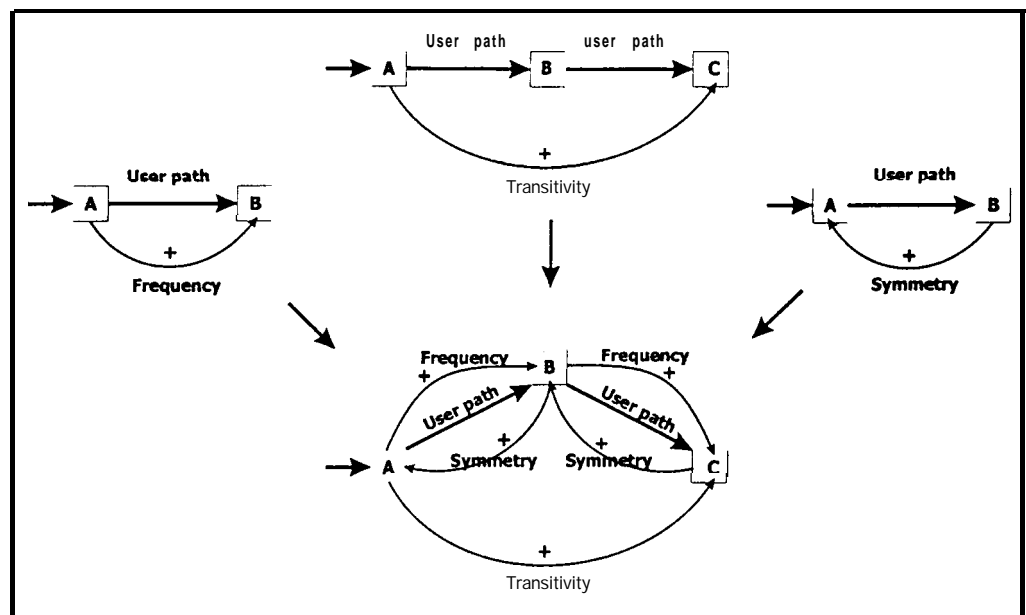


Fig. 2.: *Schematic overview of learning rules function.*

## 2.4. Link Ordering.

Changes in network structure can be fed back to users by link ordering according to descending connection strength. Any page in the network would contain an ordered list of links in which the strongest connections appear at the top. The principle of ordered presentation of choice items has a number of advantages. At least one thorough analysis of collective browsing behavior has found a strong relation between the ordered





position of hypertext links and their probability of being selected<sup>14</sup>. If we can assume the system functions as intended, the most appropriate connections from a given page have the highest connection strength. They should thus appear on top of the link list to have the highest probability of being selected. Efficient link ordering has also been shown to improve selection times and reduce cognitive overhead<sup>3, 33, 38</sup>. In real hypertext systems connection strength can be communicated in many other ways such as link coloring, font size, etc, but link ordering is plausibly one of the more important factors.

### **3. Experiments with the adaptive hypertext system.**

#### *3.1. Experimental goals.*

We have tested the outlined system for adaptive hypertext restructuring in two WWW experiments to study the effects of collective browsing in an adaptive hypertext system on the WWW.

Rather than implementing a finished system for real hypertext networks we adopted a more experimental approach that would enable us to gather reliable and valid data on the system's basic functioning. Although the face validity of our results would be reduced, this would provide us a firm, empirical validation of the system's main principles before applying it to actual hypertext systems with all subsequent biases. Our approach led to the reduction of our system to one that implemented hypertext nodes in a very reduced form (words and hyperlinks alone) and instructions to participants to use one, specific browsing strategy.

Systems that dynamically restructure hypertext networks according to the implicitly measured preferences of their users are subject to a number of biases. First, the content of the specific hypertext used can in itself induce a specific hypertext network structure. A network containing high level nodes such as "mammal", "fish", "insect" combined with less abstract nodes on actual animals like "cats", "salmon" and "fly" could induce the network's users to select links according to the hierarchical nature of the network's content. The resulting network structure would thus be shaped by the experimenter's selective choice of nodes as well as the adaptive hypertext system itself. Both effects would be difficult to separate and the specific *success* or failure of our approach would be difficult to assess from our results. This bias can not be entirely avoided because a hypertext network needs actual nodes, but it can be reduced to a maximal extent by reducing the nodes' content and applying a strict a-select (i.e. unbiased selection) and random criterion for the selection of network nodes.

Second, users are known to use a myriad of personal information seeking strategies<sup>25</sup>, depending on retrieval goals, context, past experiences, etc. Clearly, the development of hypertext networks that

shape their structure according to user selections could indirectly be influenced by the distribution of specific browsing strategies among the population of its users. This distribution is unknown and can be affected by artifacts of the system's set-up and content. Instructing participants to apply one, specific browsing strategy, preferably one that encourages subjects to randomly and unbiasedly browse the network, could possibly control this bias. With a sufficiently large group of participants individual, random biases have a higher probability of being cancelled out.

### 3.2. Set-up

#### 3.2.1. Network Nodes.

The 150 most frequent English nouns were derived from the LOB-corpus<sup>17</sup> and used as nodes for the experimental network (see table 1). This set of nouns included abstract words like "influence", "time", "system" and "development" as well as less abstract words like "father", "water", "car" and "building". 150 seemed to be a reasonably large number of nodes to provide users a rich and large enough network to browse and to train, while the resulting data would still be manageable in later analysis.

Our use of the 150 most frequent English nouns can be justified by the assumption that frequent nouns have a clear and well understood meaning among speakers of the language. Most importantly, however, word frequency is a relatively a-select criterion for the selection of words.

*act, action, age, amount, area, art, attention, authority, bed, blood, board, body, book, boy, building, car, case, century, change, church, city, committee, commonwealth, company, conference, control, council, country, course, court, day, death, development, door, doubt, education, effect, end, evening, evidence, example, experience, face, fact, family, father, field, figure, film, food, form, friend, girl, government, group, hand, head, health, heart, history, house, idea, industry, influence, interest, job, kind, knowledge, labour, land, language, law, level, life, light, line, love, man, market, matter, meeting, method, mind, moment, money, morning, mother, movement, music, name, nature, need, night, number, office, order, paper, part, party, peace, period, person, place, point, policy, position, power, problem, question, rate, reason, research, result, road, room, school, section, sense, service, side, situation, society, sort, stage, state, story, system, table, tax, theory, thing, thought, time, town, trade, training, type, use, value, view, voice, water, way, week, wife, woman, word, work, world, year*

Table 1: *The 150 most frequent nouns in the English language according to the L.O.B. corpus in alphabetical order.*

#### 3.2.2. Software and interface.

A HyperCard application was set up as a CGI-script<sup>11</sup> and system management tool for each experiment. The application contained the network nodes, their weighted connections to all other nodes in the



network and applied the system's learning rules. At the start of each experiment, the HyperCard application initialized the connections between all words in the network to small random values ( $<0.1$ ) for a first random ordering of hyperlinks.

Upon first contact the program assigned each participant a random starting position in the network from which the browsing session could start. The application would then generate the appropriate hypertext pages from which the participants could browse the network. The generated hypertext pages consisted of no more than:

- 1) a header denoting the user's current position in the network
- 2) a list of the 10 highest ordered hyperlinks from that position
- 3) and a "more items"-link that users could select to see the next 10 ordered links from the list and so on until the 10 last ordered ones.

(The rationale for the reduction of our network nodes to this minimal form has been explained in section 3.1.) When the user selected one of the listed hyperlinks, the HyperCard stack was contacted via our WWW server and shifted the user's position to the selected node by returning a new hypertext page corresponding to the requested one with a new ordered list of links. At each consequent selection the adaptive system would apply the learning rules to the selected connections.

The hyperlinks in the pages contained so-called packed queries<sup>11</sup>, i.e. their URL's (Universal Resource Locator: a **website** address) were generated in advance to contain the names of the two, last requested nodes. When users selected a link from the list, the URL itself thus informed the adaptive system of where the user was coming from so it could execute the transitivity and symmetry learning rules.

### 3.2.3. Learning rules application.

The HyperCard program was set up to apply two learning rules in the first in situ experiment, i.e. Frequency and Transitivity, and to extend these with the Symmetry learning rule for the second experiment. Each learning rule applied a different reward to the connections they reinforced. The rewards were independent from the connection's weight.

First, the frequency learning rule was set to apply a reward of 1, so each time a user traversed the connection between a node *a* and a node *b* this connection's weight was increased by a value of 1.

Second, the transitivity learning rule was set to apply a reward of 0.5. Each time a user used a hyperlink between a node *a* and a node *b*, and consequently the hyperlink between node *b* and a node *c*, the connection between node *a* and node *c* was increased by a value of 0.5. The reward for the Transitivity learning rule was set to this level because the rule, contrary to the Frequency learning rule, rewarded connections that hadn't actually been selected by the user. Although transitive

connections are highly plausible in many cases, they can be spurious. For example, if a user connects cat and milk, and afterwards connects milk with cheese, the transitive connection, i.e. cat and cheese is not necessarily a valid one. The lower value of 0.5 was chosen to reflect the estimated probability of rewarding these spurious links.

Thirdly, the Symmetry learning rule, which was only used in the second experiment, applied a reward of 0.3. So, each time a user used the connection between a node *a* and a node *b*, the connection from *b* to *a* was, increased with a reward of 0.3. The value of this learning rule's rewards reflected our estimation of how relevant symmetric connections would be compared to the transitive ones.

### 3.3. *Subjects and instructions.*

After setting up the adaptive hypertext system we sent requests to participate to all relevant mailing lists and newsgroups. The request included a sketchy overview of what the experiment was like and a URL that pointed to the WWW server on which the experiment was conducted. Participation in the experiment was entirely voluntary and anonymous. We registered no personal information but participants were encouraged to answer a few personal questions after having completed the experiment.

When contacted our server returned a brief explanation of the experiment's interface and general instructions concerning browsing strategies, etc. No mention was made of the experiment's goal or implementation. In general we advised people to browse the network as they would browse any other hypertext network<sup>5,6,7</sup>. Rather than feeling tested and expressing the most personal associations and links, subjects should associatively wander from page to page and choose hypertext links that were most related to their present position in the network. (The reasons for this reduction of browsing strategies are outlined in section 3.1.) The one-page instructions ended in a hypertext link that pointed to the actual experimental network. User could exit the network at any time by selecting "exit".

### 3.4. *Results.*

#### 3.4.1. Participation

A total of 12,000 requests were received in both experiments. The first hypertext experiment received as many as 4,600 requests while the second experiment received at least 6,000 requests. Samples from the server logs showed that participants individually selected about 10 links per session. If most people tried the experiment once and did not return, then the two experiments attracted an estimated total of 600 participants each. Both experiments lasted about one month. The experimental

hypertext system saved a copy of network structure every 200 requests for later analysis.

### 3.4.2. Link development.

Both experiments started from an entirely random structure, but rapidly developed meaningful connections from all nodes. Although users were set to start browsing from random positions in the network, the browsing activity seemed to concentrate on a rather limited set of nodes in the network. This effect seemed to be more pronounced in the first experiment. Nevertheless, most nodes in both experiments were reasonably frequently retrieved and browsed to develop a set of relevant connections to the other nodes in the network. An example might demonstrate how links developed in the network. Table 2 shows how connections from the node MIND developed during the first 4,250 established user selections in the second experiment. At first the ordered list of links from the node MIND consisted of nothing more than random, meaningless connections as a results of the initial assignment of random connection weights. After 600 selections a number of meaningful connections like 'Thought', 'Idea' and 'Research' showed up on top of the list. A number of lower ordered connections such as 'Law' and 'Light' however remained, as remnants of the initial random ordering. After about 1,200 links, all ordered links seemed to be sufficiently associated with MIND, but the order of links in the list was still being refined. 'Knowledge' and 'Development' for example had shifted upwards and had passed 'Research' that had shifted down. After about 2,400 connections the list with connections seemed to have stabilized as no new connections appeared. Only the order of the links changed slightly. The list of links from MIND had achieved a more or less stable structure in which each connection had taken a position that best reflected its semantic or associative strength of relation to MIND.


MIND					
Connections:					
1st	Table	Thought	Thought	Thought	Thought
2nd	Order	Idea	Idea	Idea	Idea
3rd	Figure	Research	Knowledge	Knowledge	Knowledge
4th	Party	Problem	Change	Theory	View
5th	Question	Need	Development	Development	Education
6th	School	Light	Research	Change	Theory
7th	Act	Development	View	Research	Development
8th	History	Change	Need	View	Research
9th	Fact	View	Problem	Need	Change
10th	Wife	Law	Theory	Education	problem
# Training (links):	0	600	1200	2400	4200

Table 2: Development of connections from the network node MIND from initial random state to condition after 4,200 selections

### 3.4.3. Development speed for both experiments.

The symmetry learning rule was added in the second experiment to speed up network development by allowing the network to produce a greater variety of links that could be selected upon by the Frequency learning rule. How did this influence network development speed?

Network development speed can be operationalized as the difference in QAP network correlation<sup>15</sup> between each consequent network structure and the first registered network structure. If the difference in correlation between two consequent stages of network development and the first network is low, little development has taken place during the two consequent measures. If, on the other hand, this difference in correlation is high, the speed of network development has been high during the two measurements. Since the networks start from a random structure the second first network structure is correlated to all consequent networks. Graph 1 shows network development for both experiments. The graph shows that network development in both networks is very fast during the first 1,800 connections and then decelerates. Network development in the second experiment is slightly faster than in the first experiment for the first phases of network development, but reaches its asymptote sooner afterwards.

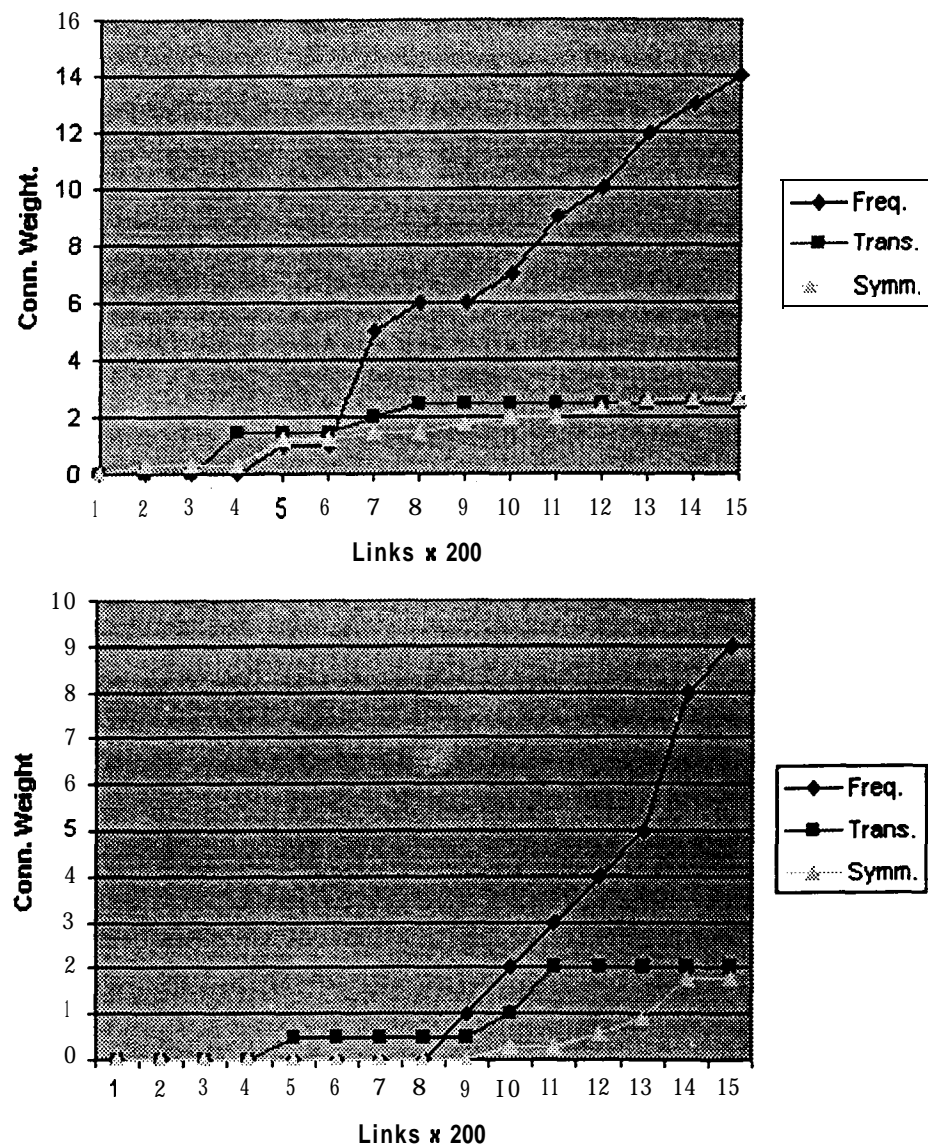


Graph 1: **Network development expressed in QAP correlations with first network as a function of the number of selections in the network.**

### 3.5. Learning rules interaction.

All learning rules operate in parallel and are expected to interact in the creation of new links and their consequent selection and reward by the user's choices. To analyze the learning rules' different contributions, we

plotted the learning rule's rewards to the 20 strongest connections in the network from the second experiment. Six of these 20 connections have first been awarded by the Symmetry or Transitivity learning rule, then being picked up by repeated rewards from the Frequency learning. Due to the rather large interval of connections between subsequent measurements of network state, apparently 'simultaneous' onsets of rewards from all three learning rules could not be resolved into separately measured events, and so we expect this number to be even higher in reality. Graph 2 shows the rewards for the connections between the nodes Knowledge and Research and the connection between the Life and Nature nodes, as a function of the number of links that had been established in the network. Both connections have been introduced by either the Symmetry or Transitivity learning rule.



Graph 2: Contributions of different learning rules rewards for Knowledge-Research and Life-Nature connections as a function of number of selections in the network.

### 3.4. Network structure.

We performed an n-clique cluster analysis on the results from both experiments to gain insight into the general network structure. The cluster analysis revealed a remarkable structure of closely connected sets of nodes that seemed to be grouped according to their relation to highly general or abstract concepts such as “Time”, “Space” and “Cognition”. Table 3 provides an overview of the nine clusters that were found in the network structure resulting from the second experiment. Both networks contained highly similar clusters. In both cases the “cognition” cluster takes up about 33% of all words over all clusters, indicating its central position and importance in the network.

Cluster name	Cluster Content
Time	Age, time, century, day, evening, moment, period, week, year
Space	Place, area, point, stage
Movement	Action, change, movement, road, car
Control	Authority, control, power, influence
Cognition	Knowledge, fact, idea, thought, interest, book, course, development, doubt, education, example, experience, language, mind, name, word, problem, question, reason, research, result, school, side, situation, story, theory, training, use, voice
Intimacy	Love, family, house, peace, father, friend, girl, hand, body, face, head, figure, heart, church, kind, mother, woman, music, bed, wife
Vitality	Boy, man, wife, health
Society	Society, state, town, commonwealth
Office	Building, office, work, room

Table 3: Results of n-clique cluster analysis on networking resulting from second experiment. Cluster names express intuitive descriptions of cluster content

We also calculated the QAP-correlation between the two final networks to see how the differently developed network's structure corresponded. The calculated QAP correlation measured 0.58 after 4,200 jumps. When symmetric closure was added to the first network's structure the QAP correlation measured 0.63. Both experiments thus resulted in reasonably similar networks.

### 3.7. Discussion.

#### 3.7.1. Network development and feedback.

These results indicate that network development was very fast and occurred mostly during the first phase of the network's development. In their first 1,800 links both networks quickly achieved a relatively stable set of connections from most nodes. Afterwards only minor changes to the network structure took place in the form of small adjustments to the different ordering of links and connection weights.

The pattern in which a fast and vigorous development of network structure occurs during the first phases of the training process to then slow down to reach an asymptote can be explained by two factors. First, the rewards from the learning rules were set at constant values that were



simply added to the connection's actual value. Connection strengths thus grew linearly larger with an increased number of selections and the rewards from the learning rules thus exerted continually less influence on the ordering of links in the network during the last phase of network development. This effect can account for the slow rate of change in network development after a certain number of selections. Second, the rapid restructuring of network structure during the initial phases of network development can be explained by a positive feedback loop that might be involved in the above described experimental set-up, and in fact the entire concept of a self-organizing network. As could be expected, subjects are more likely to select the items they read first in the list proposed to them. Therefore, connections that rise in the rank ordering because they are selected and rewarded would have a significantly higher probability of being selected on a following occasion. Thus, reinforcement of a link tends to produce further reinforcement. This feedback loop between link ordering and link selection might cause the high speed of network development during the first 1,800 links. Feedback can strengthen new connections introduced by Symmetry and Transitivity in a fast loop of continuing rewards, administered by the Frequency learning rule. Any worthwhile link could as such rapidly achieve a high position, and pull transitively or symmetrically related links up in its wake. Our analysis of the different contributions of each learning rule to the 20 strongest connections in the second network (graph 2) and the analysis of network development speed (graph 1) seem to confirm this assumption.

### 3.7.2. Cluster analysis and semantic attractors.

At the end of each experiment, after some 6,000 selections, the most frequented nodes had gathered a list of 10 strongest links that quite well reflected their direct semantic environment, with words that were near synonyms of the node name at the top of the list. However, this positive result was much less strong in the less frequented nodes, because of what we termed the "attractor effect". Nodes that had many incoming links, by accident, or because they were associated with many other words in the list, would tend to attract more users. This would result in increasing strength of their incoming paths, and their replacement by even stronger direct links. Especially in the first experiment, almost all paths would end up in a cluster of semantically related, strongly cross-linked nodes, forming an approximate attractor for the network (cf. 'cognition'-cluster in table 3). Although the random assignment of starting nodes meant that all nodes would be consulted on first entry with the same average frequency, the subsequent moves would very quickly end up in the attractor cluster. As a result, nodes outside the attractor would get little chance to learn and thus remained poorly connected.

In our second experiment, the introduction of the Symmetry rule attenuated this effect, since strong links leading into an attractor would necessarily produce weaker, inverse links leading out of the attractor. This gave nodes outside the attractor the chance to develop some links of their own, generating new local attracting clusters, weakly connected to other clusters. The overall learning seemed more efficient in the sense that less time was needed to develop good associations, and the result was more balanced, in the sense that the differences in frequentation between nodes were less strong. The data shown in graph 1 as well as the cluster analysis of the networks seems to support this thesis. Although the learning algorithms only work on links and not on groups of nodes, it is remarkable how well the resulting clusters fit in with intuitive categories. This also seems to confirm that the set-up achieves its aim of absorbing the common semantics of a heterogeneous group of users.

#### **4. Experiments with the artificial navigator.**

##### *4.1. Introduction.*

The question however remained to what extent this system of adaptive hypertext succeeded in restructuring itself to ensure a maximal resemblance to the users' actual shared link preferences. The data from our first and second in situ experiment lacked any direct **CROSS-**measurement of the users' preferences and we were thus not able to compare these to the final network structure. We therefore performed a number of simulations in which a programmed navigator browsed and trained the adaptive hypertext networks rather than human subjects. This would enable us to train a sufficient number of networks to measure the statistical resemblance between the trained networks and the artificial navigator's known preferences. To enable comparisons with the previous in situ experiments, we used the same networks and learning rules. The only difference was that one and the same artificial navigator, whose mental model was perfectly accessible and stable, would browse these networks.

##### *4.2. A heuristic model of browsing in hypertext.*

To implement the artificial navigator, we needed to plausibly simulate the behavior of human navigators in hypertext. Therefore, the following plausible and computationally viable model of hypertext browsing was developed <sup>3</sup>. It is based on the outlined analyses of the hypertext-user system and the heuristic procedure known as hill-climbing <sup>32</sup>. Hill climbing is a search heuristic that tries to locate a given function's global optimal values by always shifting position towards locally optimal values.

Let us assume that we can measure both network structure and users' preferences and represent these by matrices.

Then,

$U$ , is the square, Boolean matrix of order  $n$  representing the user model (hypertext network structure),  $n$  being the number of pages in the network

$M$ , is the square, integer matrix of order  $n$  that represents a subset of the user's preformed common knowledge of associations in general <sup>20</sup>, applied to the nodes of network  $U$ , i.e. the users mental model.

Assume the user has started browsing at position  $s$  in the network. The hyperlinks from this position can be found in  $U$ 's row vector  $v_s$ . If the user is attempting to retrieve a certain node  $g$  in the network, then he will have to evaluate the values of vector  $v_s$  according to how their selection will lead the user closer to the desired position  $g$ . The values that shape this evaluation can be found in matrix  $M$ 's column vector  $v_g$ . This vector's values indicate the extent to which the user feels a certain hyperlink will lead closer or farther from the desired goal-position.

If the values in both vectors,  $v_g$  and  $v_s$  are multiplied pair-wise the resulting vector will indicate how useful a certain connection will be to move closer to the desired position in the network. The user will then move to the position associated with this vector's maximal value from which the same evaluation will take place, and so on, until the goal position has been achieved.

#### *4.3. Implementation and results of the artificial navigator.*

The outlined model of associative hill climbing in hypertext has been implemented in an artificial navigator agent that was used to independently browse and train adaptive hypertext networks. Lacking adequate data the artificial navigator's mental model was derived from the weighted average of the networks, from our in situ experiment. The use of our data from the previous experiments seems to introduce circularity to the design. We are however not interested in the agent's exact browsing behavior but rather the similarity between the user's known mental model and the resulting hypertext network structure. Any sufficiently structured mental model would thus suit our purposes.

The artificial navigator was programmed in C and ran on the university's IRIX system. The artificial navigator trained 20 adaptive hypertext networks that were each browsed for 3,000 connections.

The following interesting measurements could be derived from these results. Exactly the same (artificial) user using exactly the same mental model had trained the 20 initially randomized networks. Therefore any variations in the resulting network structure would be a product of the network's learning process. Two kinds of deviations could occur. First, the resulting networks could be unreliably related to the (artificial) user's mental model. Unreliability in measurement occurs



when for the same measured quantity a different measurement value occurs on different occasions. This bias is mostly due to random aberrations in the measuring apparatus. Second, even if the networks are reliably related to the (artificial) user's mental models they can still be invalid due to constant biases in the measuring apparatus. The measurement is then said to be reliable but invalid; it measures constant, but wrong values.

Both reliability and validity of the adaptive hypertext system can be calculated from the networks that have been trained by the artificial navigator. The amount to which the trained networks correlate among each other is an indication of the stability of network development. Validity of network development on the other hand can be expressed by the correlation between the artificial navigator's known mental model and the final network structure.

The following QAP-correlations have thus been calculated.

- 1) the QAP correlations between the resulting hypertext networks as an indication of the reliability of network development averaged to 0.79, indicating a high reliability of network development
- 2) The QAP correlation between the user's mental model and the resulting hypertext network structure averaged to 0.82, indicating a high validity of measurement

Both calculated QAP-correlations thus indicate that development of hypertext networks is reliably and validly capable of resembling the simulated user's mental model.

## **5. General Conclusion.**

The presented system for adaptive hypertext structuring seems to provide us with one more way of dealing with the difficulties of designing adequate hypertext networks. Rather than using separately measured and maintained user models, the system attempts to adaptively structure hypertext networks themselves into valid and reliable models of user group preferences. It does so without explicit relevance feedback and is based on very simple and plausible assumptions about the preferences underlying users' browsing behavior. Since the network is shaped into a model of its users' overlapping preferences, the network itself functions as collaborative filtering mechanism: it "advises" users to preferentially choose one hyperlink over another by ordering hyperlinks according to connection weight.

The system has been shown to be able to structure random networks of reduced hypertext pages into well-structured networks. Simulations indicate that the resulting network adequately correspond to the users' link preferences. Network development is reasonably fast and does not seem to require an excessive amount of training.

The 'system however has been experimentally reduced to only a shallow representation of what actual hypertext systems are like. Nodes were reduced to simple word headers and an ordered list of hyperlinks. Hyperlinks were presumed to indicate exactly what they were referring to. Users were instructed to use only one of a wide variety of possible browsing strategies, namely random, associative browsing. We felt this reduction was necessary to test the system's unbiased potential, but it did however strongly reduce the face validity of our results.

One other disadvantage of the present approach is that hypertext networks develop as representations of group link preferences. Users that have different preferences from the group training the network will be faced with a network structured according to the group's overall preferences. This structure will offer most users within the group the best chances of finding the link pattern they desire, but they do not reflect personal, deviating preferences. Individuals can, however, train small networks and these networks will reflect these individuals' personal link preferences.

In spite of these shortcomings, we believe the system in its present form can be applied to real hypertext networks. First, the system could function as a link recommendation system on top of an existing hypertext network. The present, ordered list of connections could then be implemented as a bar underneath actual hypertext pages, much like the **Alexa** system does at present. While leaving the existing hypertext network structure intact, this would offer users the added advantage of adaptive hyperlinking. After the adaptive links have settled in a more or less stable state, they could inform hypertext designers about the users' preferences so that they could make the necessary adjustments to existing link structure.

Second, we envision a system using Java Applets that could replace hyperlinks in existing hypertext networks. Rather than having a fixed part of the text (i.e. the anchor) refer to a fixed and single hypertext document, a Java Applet could replace the static anchor with an ordered lists of recommended hyperlinks (including the already existing one). The list of hyperlinks could be generated by a localized version of the outlined adaptive hypertext system.

Third, since the reliability and validity of the outlined system for adaptive hypertext have been demonstrated in the artificial navigator simulations, the system could be applied to measure user link preferences in their own right. Public websites could for example make use of the system to measure its audience's associative preferences. The **websites** could be set up to contain the specific concepts of interest as nodes (e.g. "soap powder", "whiteness" and "cleanness"). While users browse the web site, they shape its structure into a representation of how they feel

the concepts or nodes of that website should relate. The resulting network structure could then be used as a valid measurement of the public's "mental map" of these concepts. This approach is being studied now and is expected to deliver its first applications for marketing research within the next year.

Future research will focus on issues concerning the validity of our approach and extensions to the present learning mechanism. The experiments did not for example include independent measurements of its participants' associative preferences. The artificial navigator simulations demonstrated the learning system's ability to validly and reliably reflect its users combined ideas of the preferred relations among nodes, but its results depend entirely upon the validity of its model of browsing behavior. In order to compare the resulting network's structure with alternative measures of the participants' preferences, new experiments will be conducted in which subjects before entering the experiment will have to participate in a prior, explicit measurement of word association norms for concepts in the network. This data could then later be compared to the final network structure.

Likewise, new experiments will have to focus on how different values for the learning rules parameters affect learning efficiency and whether instructing users to apply different browsing strategies will affect our results. We do not expect changes in learning parameters or browsing strategies to significantly affect the development of network structure for the following reasons. First, the learning rules' parameters have already been manipulated across the two experiments. The symmetry learning rule could be said to have applied a zero reward in the first experiment and a small 0.3 reward in the second. The addition of this entirely new learning rule did not significantly affect the two resulting network structures. However, when taken to extremes the learning rule's rewards will affect our results and as we lack information about these critical values, further experimentation is warranted. Second, although users might apply different, more goal-oriented browsing strategies when training the network, this will not necessarily affect the experiment's main outcome. The presented learning mechanism uses strictly local information based on the simplest assumption that users frequently select valuable links. All three learning rules operate within a domain of maximally two consequent connections and are thereby largely insensitive to users' long term goals and behavior in the network. Browsing strategies that encompass decisions spanning over more than two consequent connections will only slightly influence learning. Further research will indicate to what extent this is the case.

One other issue that needs to be addressed is the appropriateness of our approach to WWW applications. Although the system has been



designed for the WWW (distributed learning, localized, implicit measurement of user preferences) and tested via the WWW, it can not be applied to WWW hypertext restructuring as it is today. An implementation on the web would require a change to the HTTP protocol that enables servers to inform each other about where users are coming from and which connections to reward. Presently, Java Servlets and Remote Method Invocation techniques could possibly offer an interesting way to tackle this problem.

We believe that although the present system has its limitations and should be extended with more complex algorithms and ways of dealing with the complexity of real hypertext complexity, its simplicity and generality will contribute to the domain of user modeling and automated hypertext structuring.

## References

1. ALLEN, R. B. Mental Models and User Models. In: M. Helander, T.K. Landauer & P. Prabhu, eds. **Handbook of Human-Computer Interaction**. Amsterdam: Elsevier, 1997.
2. BERNERS-LEE, T., CAILLEAU, R., GROFF, J.F. & POLLERMAN, B. , World Wide Web: The Information Universe. **Electronic Networking: Research, Applications and Policy**, 1(2), 1992, 74-82.
3. BOLLEN, J. Cognitive Complexity vs. Connectivity: efficiency analysis of hypertext networks. In: F. Heylighen & D. Aerts, eds. **The Evolution of Complexity**. Dordrecht: Kluwer Academic Publishers, 1998. [in press]
4. BRUSILOVSKY, P. Methods and techniques of adaptive hypermedia. **User Modeling and User-Adapted Interaction**, 6(2-3), 1996, 87-129.
5. CANTER, D., RIVERS, R. & STORRS, G. Characterising User Navigation through Complex Data Structures. **Behavior and Information Technology**, 4(2), 1985,93-102.
6. CATLEDGE, L. D. & PITKOW, J. E. Characterizing Browsing Strategies in the World Wide Web. **Computer Networks and ISDN Systems**, 27, 1995, 1065-1073.
7. COVE, J. F. & WALSH, B. C. On-Line text Retrieval via Browsing. **Information Processing and Management**, 24(1), 1988, 3 1-37.
8. DIENES, Z. Connectionist and Memory Array Models of Artificial Grammar Learning. **Cognitive Science**, 16, 1992, 4 1-79.
9. DILLON, A. Myths, Misconceptions, and an Alternative Perspective on Information Usage and the Electronic Medium, in: J.F. Rouet, ed. **Hypertext and Cognition**. Lawrence Erlbaum Publishers, New Jersey, 1996.
10. EDWARDS, D.M. & HARDMAN, L. (1989), Lost in Cyberspace: Cognitive Mapping and Navigation in a Hypertext Environment. In: R. McAleese, ed. **Hypertext: Theory into practice**. New Jersey: Ablex Publishing Corporation, 1989.
11. GUNDAVARAM, S. CGI programming on the World Wide Web. Bonn: O'Reilly and Associates, 1996.
12. HEBB, D. O. The organisation of behavior: a neuropsychological theory. New York: Science Editions, 1967.
13. HINTON, G. & ANDERSON, J. Parallel Models of Associative Memory. New Jersey: Lawrence Erlbaum Associates, 198 1.
14. HUBERMAN, B. A. PIROLI, P. L. PITKOW, J. E. LUKOSE, R. M. Strong Regularities in World Wide Web Surfing. **Science**, 280 (5360). 1998.95-97.

15. HUBERT, L. & SCHULTZ, J. Quadratic Assignment as a general data analysis strategy. **British Journal of Mathematical and Statistical Psychology**, **29**, 1976. 190-241.
16. JOACHIMS, T., FREITAG, D., MITCHELL, T. Webwatcher: A Tour Guide for the World Wide Web, Proceedings of IJCAI97. 1997.
17. JOHANSSON, S. & HOFLAND, K. Frequency analysis of English vocabulary and grammar: based on the LOB corpus. Oxford : Clarendon Press, 1989.
18. JONASSEN, D. H. Effects of semantically structured Hypertext Knowledge Bases. In: C. McKnight, A. Dillon and J. Richardson, eds. **Hypertext: A Psychological Perspective**. New York: Ellis Horwood, 1993.
19. JONASSEN, D.H. Semantic Network Elicitation: Tools for structuring of Hypertext. In: R. McAleese & C. Green, eds. **Hypertext: The State of the Art**. London: Intellect, 1990.
20. KAPLAN, C., FENWICK, J. & CHEN, J. Adaptive Hypertext Navigation Based on User Goals and Context. **User Models and User Adapted Interaction**, **3(2)**, 1993, 193-220.
21. KOHONEN, T. Associative Memory: a System-Theoretical Approach. Berlin: Springer-Verlag, 1978.
22. KUIPERS, B. On Representing Commonsense Knowledge. In: N. V. Findler. ed. **Associative Networks - Representation and Use of Knowledge by Computers**. London: Academic Press, 1979.
23. LIEBERMAN, H., Letizia: An agent That Assists Web Browsing. In: **Proceedings of the International Joint Conference on Artificial Intelligence**, Montreal. August 1995.
24. LUND, K., BURGESS, C. & ATCHLEY, R.A Semantic and Associative Priming in High Dimensional Semantic Space. **Proceedings of the Cognitive Science Society**, New Jersey: Lawrence Erlbaum Publishers, 1995, 660-665.
25. MARCHIONINI, G. Information Seeking in Electronic Environments, Cambridge University Press, Cambridge, 1995.
26. MATHE, N. AND CHEN, J. R. User-Centered Indexing for Adaptive Information Access. In: P. Brusilovsky, A. Kobsa and J. Vassileva ,eds. **Adaptive Hypertext and Hypermedia Systems**. Dordrecht: Kluwer Academic Publishers, 1998.
27. NIELSEN, J. The Art of Navigating through Hypertext. **Communications of the ACM**, **33(3)**, 1990, 298-310.
28. OSGOOD, C.E., SUCI, G.J. & TANNENBAUM, P.H The measurement of meaning. Urbana: University of Illinois Press, 1957.
29. PRABHU, P.V. & PRABHU, G.V. Human Error and User-Interface Design. In: M. Helander, T.K. Landauer & P. Prabhu, eds. **Handbook of Human-Computer Interaction**. Amsterdam: Elsevier, 1997.
30. ROCHA, Luis M. Categorizing Databases, Evidence Sets, and Evolutionary Constructivism, **International Journal of Human-Machine Studies**, **1998** [in press].
31. RUMELHART, D. E. & MCCLELLAND, J. Parallel Distributed Processing. Volume 1: Foundations. London: Bradford Books, MIT Press, 1986.
32. SCHILDT, H. C. The Complete Reference (AI-based problem solving). London: Osborne-McGraw Hill, 1987, 645-648.
33. SEARS, A. & SCHNEIDERMAN, B. Split Menus: Effectively Using Selection Frequency to Organize Menus. **ACM Transactions on Computer-Human Interaction**, **1(1)**, 1994, 27-51.
34. SNOWBERRY, K., PARKINSON, S. & SISSON, N. Effects of Help Fields on Navigating through Hierarchical Menu Structures. **International Journal of Man-Machine Studies**, **22**, 1985, 479-491.





35. STOTTS, P.D. & FURUTA, R. Petri-net based hypertext: document structure with browsing semantics. *ACM Transactions on Information Systems*, 7( 1), 1989, 3-29.
36. SZALAY, L. B. & DEESE, J. Subjective Meaning and Culture: an Assessment Through Word Associations. New Jersey: Lawrence Erlbaum Associates. 1978.
37. THURING, M., HANNEMANN, J. & HAAKE, J. M. Hypermedia and Cognition: Designing for Comprehension. *Communications of the ACM*, 38(8), 1995, 57-66.
38. TOMEK, I., MAURER, H. & NASSER M., Optimal presentation of links in Large Hypermedia Systems, in: *Proceedings of ED-MEDIA '93, World Conference on Educational Multimedia and Hypermedia*. Charlottesville: AACE, 1993.5 1 1-5 18.
39. VAN DYKE PARUNAK, H. Ordering the information graph. In: E. Berk & J. Devlin J., eds. *Hypertext and Hypermedia Handbook*. London: McGraw-Hill, 1991.299-325.

This paper was published in  
THE NEW REVIEW OF HYPERMEDIA & MULTIMEDIA  
Volume 4, 1998  
***Subscription price for Volume 4, 1998: £70.00/US\$130.00***  
***Published annually. Back issues available.***  
Taylor Graham Publishing, 500 Chesham House,  
150 Regent Street, London W1R 5FA, UK.